Learning Autonomous Vehicle Safety Concepts from Demonstrations

Karen Leung, Sushant Veer, Edward Schmerling, Marco Pavone

On the importance of having an independent safety evaluation module

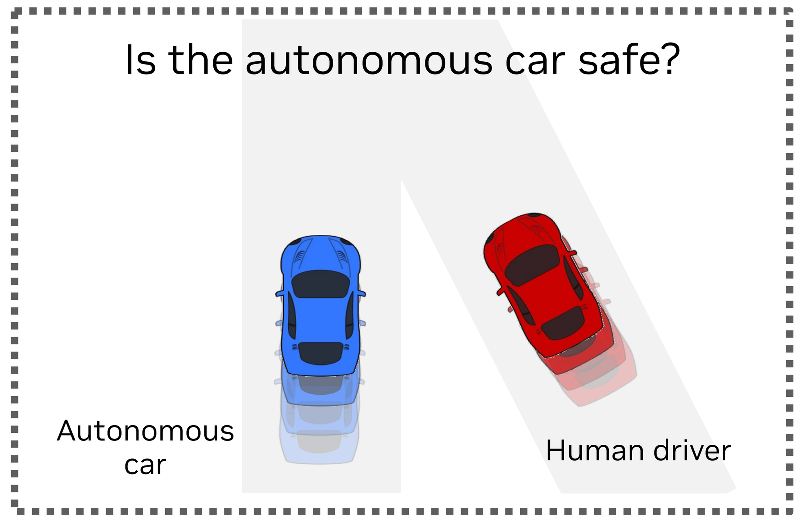

Within the AV stack, there typically is a prediction model that predicts the behaviors of other agents, and a planner that uses the prediction model to make informed decisions. However, the prediction model is not always accurate, and the planner may not respond fast enough to split-second threats.

Therefore, there needs to be an independent safety module that can intervene with a safe maneuver anytime upstream components "make a mistake". But *when* should such a safety module intervene, and *how* should it do so?

What exactly is a "Safety Concept"?

First, we define some terminology to describe what our goal is.

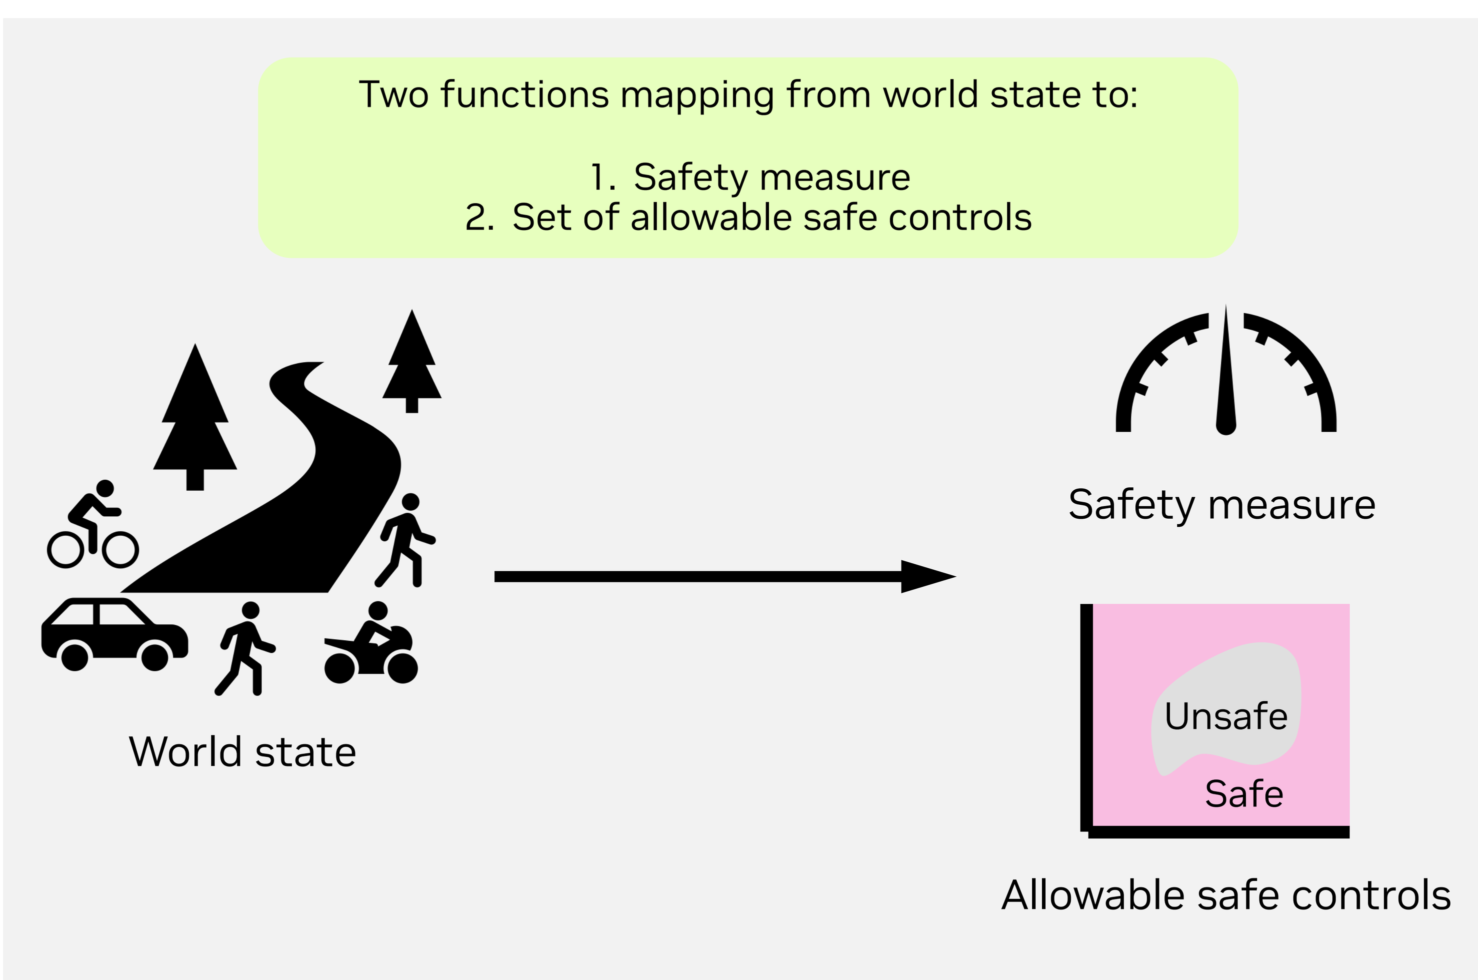

Examples:

- Velocity obstacles
- Safety Force Field
- Responsibility Sensitive Safety
- Forward reachability
- Backward reachability

How can we synthesize a novel safety concept?

We can describe a family of safety concepts via HJ reachability

Hamilton-Jacobi-Isaacs partial differential equation (Robust HJB equation)

$$\frac{\partial V(x,t)}{\partial t} + \min\left\{0, \max_{u\in U}\min_{d\in D} \frac{\partial V(x,t)}{\partial x}\cdot f(x,u,d)\right\} = 0$$
$$V(x,0) = v(x)$$

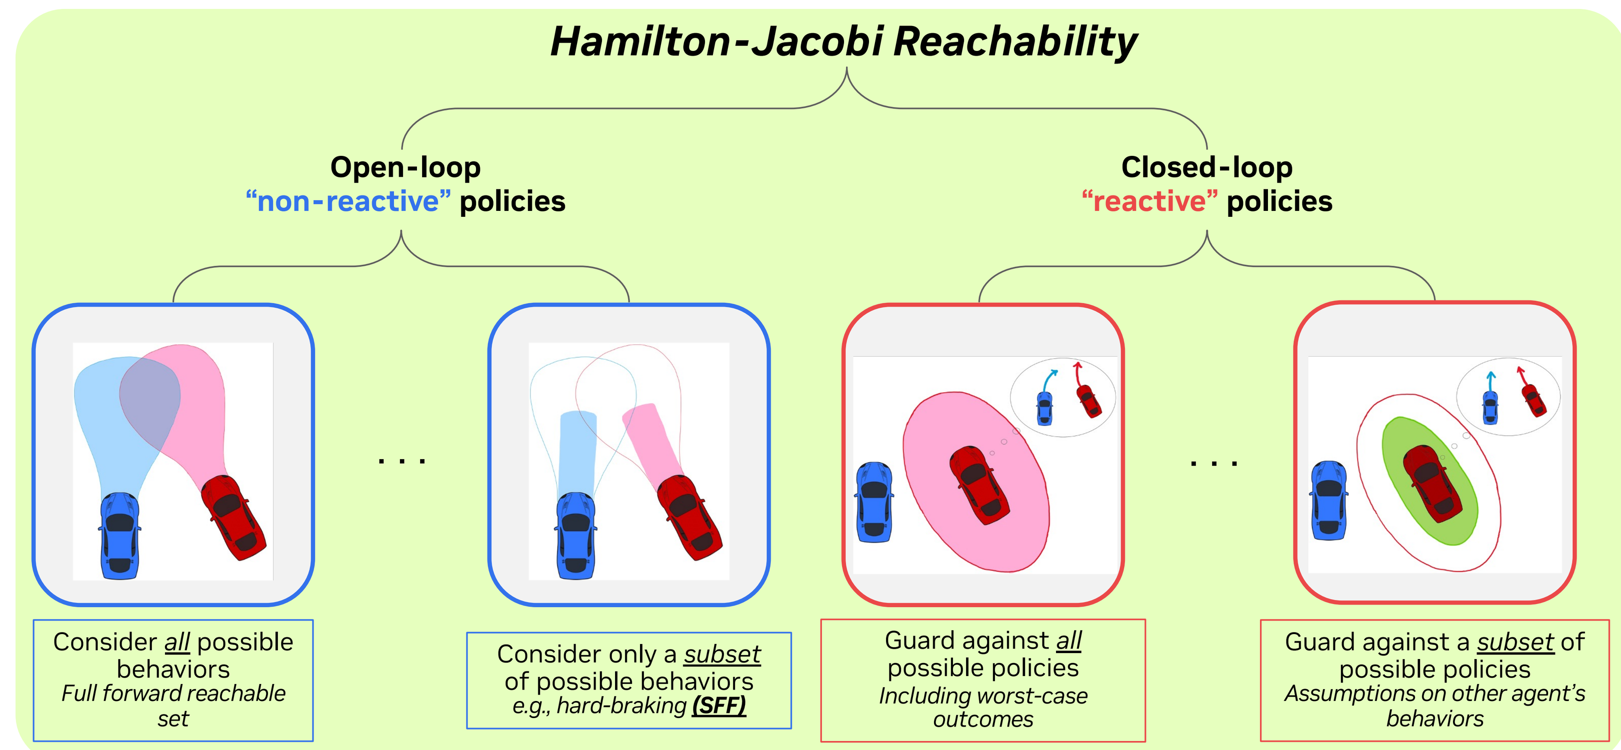

By varying the parameters of the HJ reachability problem, we can describe both closed-loop and open-loop behaviors, and anything in between.

How should we select "reasonable" control bounds?

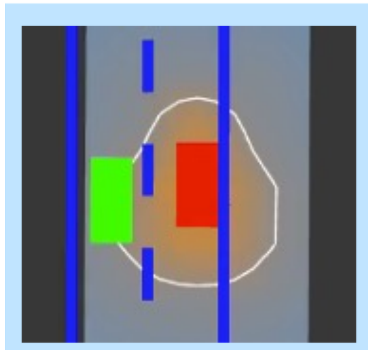

Picking *all possible* controls & disturbances leads to overly conservative safety concepts

Given a dataset of states and controls: $(x^{(1)},u^{(1)}), (x^{(2)},u^{(2)}),\ldots, (x^{(N)},u^{(N)})$ we want to learn $U(x)$

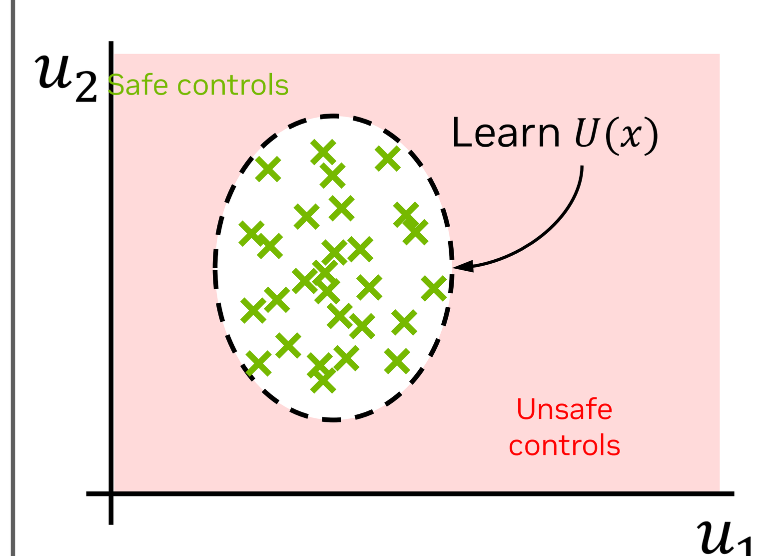

Key insight: Humans take controls that keep them safe. Taking controls outside the boundary will lead to an undesirable outcome.

Data lives inside a control invariant set.

Control set learning via Control Barrier Functions

$$\max_{u\in U} \underbrace{\nabla b(x)^T f(x,u)}_{\dot b(x)} \geq -\alpha(b(x)) \quad \forall x \in X$$

Learn parameters of α so that this condition holds for the dataset

Bound the rate at which the system approaches the boundary

How do we account for the constraint coupling when synthesizing a safety concept with HJ reachability?

Safe interaction data → CBF learning → $U(x), D(x)$ → HJ reachability → Safety concept

Constrained min-max game

Proposition 1. *Consider a coupled affine constraint* $p(u_A,u_B) := au_A + bu_B + c \geq 0$ *in* u_A *and* u_B, *and a linear objective* $q(u_A,u_B) = Au_A + Bu_B$. *Let* $\widetilde{\mathcal{U}}^A = \{u_A \in \mathcal{U}^A \mid \exists u_B \in \mathcal{U}^B, p(u_A,u_B)\geq 0\}$ *and* $\widetilde{\mathcal{U}}^B = \{u_B \in \mathcal{U}^B \mid \exists u_A \in \mathcal{U}^A, p(u_A,u_B)\geq 0\}$ *represent control sets for each agent which ensures the other agent can satisfy the constraint* $p(u_A,u_B)\geq 0$. *Let* $\mathcal{U}^A(u_B) = \{u_A \in \mathcal{U}^A \mid p(u_A,u_B)\geq 0\}$, *and* $\mathcal{U}^B(u_A) = \{u_B \in \mathcal{U}^B \mid p(u_A,u_B)\geq 0\}$ *describe an agent's feasible control set with the other agent's control fixed. Then,*

$$\max_{u_A\in\widetilde{\mathcal{U}}^A}\min_{u_B\in\mathcal{U}^B(u_A)} q(u_A,u_B) \geq \min_{u_B\in\widetilde{\mathcal{U}}^B}\max_{u_A\in\mathcal{U}^A(u_B)} q(u_A,u_B). \tag{9}$$

That is, the player that acts first has the advantage assuming the second player is provided feasible options.

What does a data-driven safety concept look like and what does it mean?